

DEEP UNSUPERVISED LEARNING THROUGH SPATIAL CONTRASTING

Elad Hoffer

Technion - Israel Institute of Technology
Haifa, Israel
ehoffer@tx.technion.ac.il

Itay Hubara

Technion - Israel Institute of Technology
Haifa, Israel
itayh@tx.technion.ac.il

Nir Ailon

Technion - Israel Institute of Technology
Haifa, Israel
nailon@cs.technion.ac.il

ABSTRACT

Convolutional networks have marked their place over the last few years as the best performing model for various visual tasks. They are, however, most suited for supervised learning from large amounts of labeled data. Previous attempts have been made to use unlabeled data to improve model performance by applying unsupervised techniques. These attempts require different architectures and training methods. In this work we present a novel approach for unsupervised training of Convolutional networks that is based on contrasting between spatial regions within images. This criterion can be employed within conventional neural networks and trained using standard techniques such as SGD and back-propagation, thus complementing supervised methods.

1 INTRODUCTION

For the past few years convolutional networks (ConvNets, CNNs) LeCun et al. (1998) have proven themselves as a successful model for vision related tasks Krizhevsky et al. (2012) Mnih et al. (2015) Pinheiro et al. (2015) Razavian et al. (2014). A convolutional network is composed of multiple convolutional and pooling layers, followed by a fully-connected affine transformations. As with other neural network models, each layer is typically followed by a non-linearity transformation such as a rectified-linear unit (ReLU).

A convolutional layer is applied by cross correlating an image with a trainable weight filter. This stems from the assumption of stationarity in natural images, which means that features learned for one local region in an image can be shared for other regions and images.

Deep learning models, including convolutional networks, are usually trained in a supervised manner, requiring large amounts of labeled data (ranging between thousands to millions of examples per-class for classification tasks) in almost all modern applications. These models are optimized a variant of stochastic-gradient-descent (SGD) over batches of images sampled from the whole training dataset and their ground truth-labels. Gradient estimation for each one of the optimized parameters is done by back propagating the objective error from the final layer towards the input. This is commonly known as "backpropagation" Rumelhart et al..

One early well known usage of unsupervised training of deep architectures was as part of a pre-training procedure used for obtaining an effective initial state of the model. The network was later fine-tuned in a supervised manner as displayed by Hinton (2007). Such unsupervised pre-training procedures were later abandoned, since they provided no apparent benefit over other initialization heuristics in more careful fully supervised training regimes. This led to the de-facto almost exclusive usage of neural networks in supervised environments.

In this work we will present a novel unsupervised learning criterion for convolutional network based on comparison of features extracted from regions within images. Our experiments indicate that by

using this criterion to pre-train networks we can improve their performance and achieve state-of-the-art results.

2 PREVIOUS WORKS

Using unsupervised methods to improve performance have been the holy grail of deep learning for the last couple of years and vast research efforts have been focused on that. We hereby give a short overview of the most popular and recent methods that tried to tackle this problem.

AutoEncoders and reconstruction loss These are probably the most popular models for unsupervised learning using neural networks, and ConvNets in particular. Autoencoders are NNs which aim to transform inputs into outputs with the least possible amount of distortion. An Autoencoder is constructed using an encoder $G(x; w_1)$ that maps an input to a hidden compressed representation, followed by a decoder $F(y; w_2)$, that maps the representation back into the input space. Mathematically, this can be written in the following general form:

$$\hat{x} = F(G(x; w_1); w_2)$$

The underlying encoder and decoder contain a set of trainable parameters that can be tied together and optimized for a predefined criterion. The encoder and decoder can have different architectures, including fully-connected neural networks, ConvNets and others. The criterion used for training is the reconstruction loss, usually the mean squared error (MSE) between the original input and its reconstruction Zeiler et al. (2010)

$$\min \|x - \hat{x}\|^2$$

This allows an efficient training procedure using the aforementioned backpropagation and SGD techniques. Over the years autoencoders gained fundamental role in unsupervised learning and many modification to the classic architecture were made. Ng (2011) regularized the latent representation to be sparse, Vincent et al. (2008) substituted the input with a noisy version thereof, requiring the model to denoise while reconstructing. Kingma et al. (2014) obtained very promising results with variational autoencoders (VAE). A variational autoencoder model inherits typical autoencoder architecture, but makes strong assumptions concerning the distribution of latent variables. They use variational approach for latent representation learning, which results in an additional loss component and specific training algorithm called Stochastic Gradient Variational Bayes (SGVB). VAE assumes that the data is generated by a directed graphical model $p(x|z)$ and require the encoder to learn an approximation $q_{w_1}(z|x)$ to the posterior distribution $p_{w_2}(z|x)$ where w_1 and w_2 denote the parameters of the encoder and decoder. The objective of the variational autoencoder in this case has the following form:

$$\mathcal{L}(w_1, w_2, x) = -D_{KL}(q_{w_1}(z|x)||p_{w_2}(z)) + \mathbb{E}_{q_{w_1}(z|x)}(\log p_{w_2}(x|z))$$

Recently, a stacked set of denoising autoencoders architectures showed promising results in both semi-supervised and unsupervised tasks. A stacked what-where autoencoder by Zhao et al. (2015) computes a set of complementary variables that enable reconstruction whenever a layer implements a many-to-one mapping. Ladder networks by Rasmus et al. (2015) - use lateral connections to allow higher levels of an autoencoder to focus on invariant abstract features by applying a layer-wise cost function.

Exemplar Networks: The unsupervised method introduced by Dosovitskiy et al. (2014) takes a different approach to this task and trains the network to discriminate between a set of pseudo-classes. Each pseudo-class is formed by applying multiple transformations to a randomly sampled image patch. The number of pseudo-classes can be as big as the size of the input samples. This criterion ensures that different input samples would be distinguished while providing robustness to the applied transformations.

Context prediction Another method for unsupervised learning by context was introduced by Dersersch et al. (2015). This method uses an auxiliary criterion of predicting the location of an image patch given another from the same image. This is done by classification to 1 of 9 possible locations.

Adversarial Generative Models: This is a recently introduced model that can be used in an unsupervised fashion Goodfellow et al. (2014). Adversarial Generative Models uses a set of networks, one trained to discriminate between data sampled from the true underlying distribution (e.g., a set of images), and a separate generative network trained to be an adversary trying to confuse the first network. By propagating the gradient through the paired networks, the model learns to generate samples that are distributed similarly to the source data. As shown by Radford et al. (2015), this model can create useful latent representations for subsequent classification tasks as demonstrated

Sampling Methods: Methods for training models to discriminate between a very large number of classes often use a *noise contrasting criterion*. In these methods, roughly speaking, the posterior probability $P(t|y_t)$ of the ground-truth target t given the model output on an input sampled from the true distribution $y_t = F(x)$ is maximized, while the probability $P(t|y_n)$ given a noise measurement $y = F(n)$ is minimized. This was successfully used in a language domain to learn unsupervised representation of words. The most noteworthy case is the word2vec model introduced by Mikolov et al. (2013). When using this setting in language applications, a natural contrasting noise is a smooth approximation of the Unigram distribution. A suitable contrasting distribution is less obvious when data points are sampled from a high dimensional continuous space, such as in the case of patches of images.

2.1 PROBLEMS WITH CURRENT APPROACHES

Only recently the potential of ConvNets in an unsupervised environment began to bear fruit, still we believe it is not fully uncovered.

The majority of unsupervised optimization criteria currently used are based on variations of reconstruction losses. One limitation of this fact is that a pixel level reconstruction is non-compliant with the idea of a discriminative objective, which is expected to be agnostic to low level information in the input. In addition, it is evident that MSE is not best suited as a measurement to compare images, for example, viewing the possibly large square-error between an image and a single pixel shifted copy of it. Another problem with recent approaches such as Rasmus et al. (2015); Zeiler et al. (2010) is their need to extensively modify the original convolutional network model. This leads to a gap between unsupervised method and the state-of-the-art, supervised, models for classification - which can hurt future attempt to reconcile them in a unified framework, and also to efficiently leverage unlabeled data with otherwise supervised regimes.

3 LEARNING BY COMPARISONS

The most common way to train NN is by defining a loss function between the target values and the network output. Learning by comparison approaches the supervised task from a different angle. The main idea is to use distance comparisons between samples to learn useful representations. For example, we consider relative and qualitative examples of the form X_1 is closer to X_2 than X_1 is to X_3 . Using a comparative measure with neural network to learn embedding space was introduced in the “Siamese network” framework by Bromley et al. (1993) and later used in the works of Chopra et al. (2005). One use for this methods is when the number of classes is too large or expected to vary over time, as in the case of face verification, where a face contained in an image has to be compared against another image of a face. This problem was recently tackled by Schroff et al. (2015) for training a convolutional network model on triplets of examples. There, one image served as an *anchor* x , and an additional pair of images served as a positive example x_+ (containing an instance of the face of the same person) together with a negative example x_- , containing a face of a different person. The training objective was on the embedded distance of the input faces, where the distance between the anchor and positive example is adjusted to be smaller by at least some constant α from the negative distance. More precisely, the loss function used in this case was defined as

$$L(x, x_+, x_-) = \max \{ \|F(x) - F(x_+)\|_2 - \|F(x) - F(x_-)\|_2 + \alpha, 0 \} \quad (1)$$

where $F(x)$ is the embedding (the output of a convolutional neural network), and α is a predefined margin constant. Another similar model used by Hoffer & Ailon (2015) with triplets comparisons for classification, where examples from the same class were trained to have a lower embedded distance than that of two images from distinct classes. This work introduced a concept of a distance ratio

loss, where the defined measure amounted to:

$$L(x, x_+, x_-) = -\log \frac{e^{-\|F(x) - F(x_+)\|_2}}{e^{-\|F(x) - F(x_+)\|_2} + e^{-\|F(x) - F(x_-)\|_2}} \quad (2)$$

This loss has a flavor of a probability of a biased coin flip. By ‘pushing’ this probability to zero, we express the objective that pairs of samples coming from distinct classes should be less similar to each other, compared to pairs of samples coming from the same class. It was shown empirical by Balntas et al. (2016) to provide better feature embeddings than the margin based distance loss 1

4 OUR CONTRIBUTION: SPATIAL CONTRASTING

One implicit assumption in convolutional networks, is that features are gradually learned hierarchically, each level in the hierarchy corresponding to a layer in the network. Each spatial location within a layer corresponds to a region in the original image. It is empirically observed that deeper layers tend to contain more ‘abstract’ information from the image. Intuitively, features describing different regions within the same image are likely to be semantically similar (e.g. different parts of an animal), and indeed the corresponding deep representations tend to be similar. Conversely, regions from two probably unrelated images (say, two images chosen at random) tend to be far from each other in the deep representation. This logic is commonly used in modern deep networks such as Szegedy et al. (2015) Lin et al. (2013) He et al. (2015), where a global average pooling is used to aggregate spatial features in the final layer used for classification.

Our suggestion is that this property, often observed as a side effect of supervised applications, can be used as a desired objective when learning deep representations in an unsupervised task. Later, the resulting representation can be used, as typically done, as a starting point or a supervised learning task. We call this idea which we formalize below *Spatial contrasting*. The spatial contrasting criterion is similar to noise contrasting estimation Gutmann & Hyvärinen (2010) Mnih & Kavukcuoglu (2013), in trying to train a model by maximizing the expected probability on desired inputs, while minimizing it on contrasting sampled measurements.

4.1 FORMULATION

We will concern ourselves with samples of images patches $\tilde{x}^{(m)}$ taken from an image x . Our convolutional network model, denoted by $F(x)$, extracts spatial features f so that $f^{(m)} = F(\tilde{x}^{(m)})$ for an image patch $\tilde{x}^{(m)}$. We wish to optimize our model such that for two features representing patches taken from the same image $\tilde{x}_i^{(1)}, \tilde{x}_i^{(2)} \in x_i$ for which $f_i^{(1)} = F(\tilde{x}_i^{(1)})$ and $f_i^{(2)} = F(\tilde{x}_i^{(2)})$, the conditional probability $P(f_i^{(1)} | f_i^{(2)})$ will be maximized.

This means that features from a patch taken from a specific image can effectively predict, under our model, features extracted from other patches in the same image. Conversely, we want our model to minimize $P(f_i | f_j)$ for i, j being two patches taken from distinct images. Following the logic presented before, we will need to sample *contrasting patch* $\tilde{x}_j^{(1)}$ from a different image x_j such that $P(f_i^{(1)} | f_i^{(2)}) > P(f_j^{(1)} | f_i^{(2)})$, where $f_j^{(1)} = F(\tilde{x}_j^{(1)})$. In order to obtain contrasting samples, we use regions from two random images in the training set. We will use a distance ratio, described earlier 2 for the supervised case, to represent the probability two feature vectors were taken from the same image. The resulting training loss for a pair of images will be defined as

$$L_{SC}(x_1, x_2) = -\log \frac{e^{-\|f_1^{(1)} - f_1^{(2)}\|_2}}{e^{-\|f_1^{(1)} - f_1^{(2)}\|_2} + e^{-\|f_1^{(1)} - f_2^{(1)}\|_2}} \quad (3)$$

Effectively minimizing a log-probability under the SoftMax measure. This formulation is portrayed in figure 4.1. Since we sample our contrasting sample from the same underlying distribution, we can evaluate this loss considering the image patch as both patch compared (anchor) and contrast symmetrically. The final loss will be the average between these estimations:

$$\hat{L}_{SC}(x_1, x_2) = \frac{1}{2} [L_{SC}(x_1, x_2) + L_{SC}(x_2, x_1)]$$

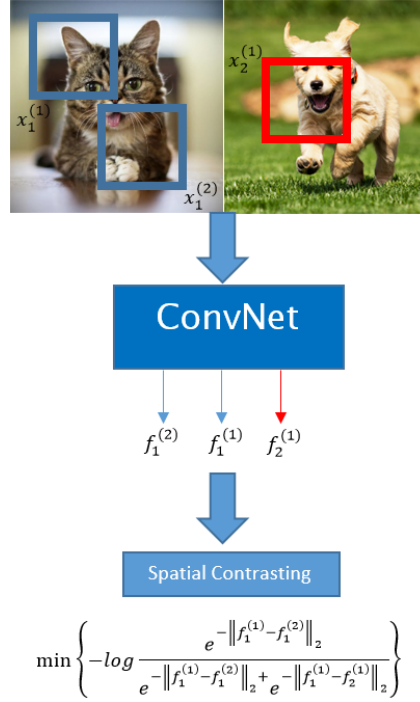


Figure 1: Spatial contrasting depiction.

4.2 METHOD

Since training convolutional network is done in batches of images, we can use the multiple samples in each batch to train our model. Each image serves as a source for both an anchor and positive patches, for which the corresponding features should be closer, and also a source for contrasting samples for all the other images in that batch. For a batch of N images, two samples from each image are taken, and N^2 different distance comparisons are made. The final loss is the average distance ratio for images in the batch:

$$L_{SC}^{-}(\{x\}_{i=1}^N) = \frac{1}{N} \sum_{i=1}^N L_{SC}(x_i, \{x\}_{j \neq i}) = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{-\|f_i^{(1)} - f_i^{(2)}\|_2}}{\sum_{j=1}^N e^{-\|f_i^{(1)} - f_j^{(2)}\|_2}} \quad (4)$$

Since the criterion is differentiable with respect to its inputs, it is fully compliant with standard methods for training convolutional network and specifically using backpropagation and gradient descent. Furthermore, SC can be applied to any layer in the network hierarchy. In fact, SC can be used at multiple layers within the same convolutional network. The spatial properties of the features means that we can also sample from feature space $\tilde{f}^{(m)} \in f$ instead of from the original image, which we use to simplify implementation. The complete algorithm for batch training is described in 1. This algorithm is also related to the batch normalization layer Ioffe & Szegedy (2015), a recent usage for batch statistics in neural networks. Spatial contrasting also uses the batch statistics, but to sample contrasting patches.

5 EXPERIMENTS

In this section we report empirical results showing that using SC loss as an unsupervised pretraining procedure can improve state-of-the-art performance on subsequent classification. We experimented with MNIST, CIFAR-10 and STL10 datasets. We used modified versions of well studied networks such as those of Lin et al. (2013) Rasmus et al. (2015). A detailed description of our architecture can be found in Appendix A.

Algorithm 1 Calculation the spatial contrasting loss

Require: $X = \{x\}_{i=1}^N$ # Training on batches of images

Get the spatial features for the whole batch of images
Size: $N \times W_f \times H_f \times C$
 $\{f\}_{i=1}^N \leftarrow \text{ConvNet}(X)$

Sample spatial features and calculate embedded distance between all pairs of images
for $i = 1$ **to** N **do**
 $\tilde{f}_i^{(1)} \leftarrow \text{sample}(f_i)$
for $j = 1$ **to** N **do**
 $\tilde{f}_j^{(2)} \leftarrow \text{sample}(f_j)$
 $\text{Dist}(i, j) \leftarrow \|\tilde{f}_i^{(1)} - \tilde{f}_j^{(2)}\|_2$
end for
end for

Calculate log SoftMax normalized distances
 $d_i \leftarrow -\log \frac{e^{-\text{Dist}(i, i)}}{\sum_{k=1}^N e^{-\text{Dist}(i, k)}}$

Spatial contrasting loss is the mean of distance ratios
return $\frac{1}{N} \sum_{i=1}^N d_i$

In each one of the experiments, we used the spatial contrasting criterion to train the network on the unlabeled images. Training was done by using SGD with an initial learning rate of 0.1 that was decreased by a factor of 10 whenever the measured loss stopped decreasing. After convergence, we used the trained model as an initialization for a supervised training on the complete labeled dataset. The supervised training was done following the same regime, only starting with a lower initial learning rate of 0.01. We used mild data augmentations, such as small translations and horizontal mirroring.

The datasets we used are:

- STL10 (Coates et al. (2011)). This dataset consists of 100,000 96×96 colored, unlabeled images, together with another set of 5,000 labeled training images and 8,000 test images. The label space consists of 10 object classes.
- Cifar10 (Krizhevsky & Hinton (2009)). The well known CIFAR-10 is an image classification benchmark dataset containing 50,000 training images and 10,000 test images. The image sizes 32×32 pixels, with color. The classes are airplanes, automobiles, birds, cats, deer, dogs, frogs, horses, ships and trucks.
- MNIST (LeCun et al. (1998)). The MNIST database of handwritten digits is one of the most studied dataset benchmark for image classification. The dataset contains 60,000 examples of handwritten digits from 0 to 9 for training and 10,000 additional examples for testing. Each sample is a 28×28 pixel gray level image.

5.1 RESULTS ON STL10

Since STL10 is comprised of mostly unlabeled data, it is the most suitable to highlight the benefits of the spatial contrasting criterion. The initial training was unsupervised, as described earlier, using the entire set of 105,000 samples (union of the original unlabeled set and labeled training set). The representation outputted by the training, was used to initialize supervised training on the 5,000 labeled images. Evaluation was done on a separate test set of 8,000 samples. Comparing with state of the art results 1 we see an improvement of 7% in test accuracy over the best model by Zhao et al. (2015), setting the SC as best model at 81.3% test classification accuracy. We also compare with the same network, but without SC initialization, which achieves a lower classification of 72.6%. This is an indication that indeed SC managed to leverage unlabeled examples to provide a better initialization point for the supervised model.

Table 1: State of the art results on STL-10 dataset

Model	STL-10 test accuracy
Zero-bias Convnets - Paine et al. (2014)	70.2%
Triplet network - Hoffer & Ailon (2015)	70.7%
Exemplar Convnets - Dosovitskiy et al. (2014)	72.8%
Target Coding - Yang et al. (2015)	73.15%
Stacked what-where AE - Zhao et al. (2015)	74.33%
Spatial contrasting initialization (this work)	$81.34\% \pm 0.1$
The same model without initialization	$72.6\% \pm 0.1$

Table 2: State of the art results on Cifar10 dataset with only 4000 labeled samples

Model	Cifar10 (400 per class) test accuracy
Convolutional K-means Network - Coates & Ng (2012)	70.7%
View-Invariant K-means - Hui (2013)	72.6%
DCGAN - Radford et al. (2015)	73.8%
Exemplar Convnets - Dosovitskiy et al. (2014)	76.6%
Ladder networks - Rasmus et al. (2015)	79.6%
Spatial contrasting initialization (this work)	$79.2\% \pm 0.3$
The same model without initialization	$72.4\% \pm 0.1$

5.2 RESULTS ON CIFAR10

For Cifar10, we used a previously used setting Coates & Ng (2012) Hui (2013) Dosovitskiy et al. (2014) to test a model’s ability to learn from unlabeled images. In this setting, only 4,000 samples from the available 50,000 are used with their label annotation, but the entire dataset is used for unsupervised learning. The final test accuracy is measured on the entire 10,000 test set.

In our experiments, we trained our model using SC criterion on the entire dataset, and then used only 400 labeled samples per class (for a total of 4000) in a supervised regime over the initialized network. The results are compared with previous efforts in table 2. Using the SC criterion allowed an improvement of 6.8% over a non-initialized model, and achieved a final test accuracy of 79.2%. This is a competitive result with current state-of-the-art model of Rasmus et al. (2015).

5.3 RESULTS ON MNIST

The MNIST dataset is very different in nature from the Cifar10 and STL10, we experimented earlier. The biggest difference, relevant to this work, is that spatial regions sampled from MNIST images usually provide very little, or no information. Because of this fact, SC is much less suited for use with MNIST, and was conjured to have little benefit. We still, however, experimented with initializing a model with SC criterion and continuing with a fully-supervised regime over all labeled examples. We found again that this provided benefit over training the same network without pre-initialization, improving results from 0.63% to 0.34% error on test set. The results, compared with previous attempts are included in 3.

6 CONCLUSIONS AND FUTURE WORK

In this work we presented spatial contrasting - a novel unsupervised criterion for training convolutional networks on unlabeled data. Its is based on comparison between spatial features sampled from a number of images. We’ve shown empirically that using spatial contrasting as a pretraining technique to initialize a ConvNet, can improve its performance on a subsequent supervised training. In cases where a lot of unlabeled data is available, such as the STL10 dataset, this translates to state-of-the-art classification accuracy in the final model.

Table 3: results on MNIST dataset

Model	MNIST test error
Stacked what-where AE - Zhao et al. (2015)	0.71%
Triplet network - Hoffer & Ailon (2015)	0.56%
Jarrett et al. (2009)	0.53%
Ladder networks - Rasmus et al. (2015)	0.36%
DropConnect - Wan et al. (2013)	0.21%
Spatial contrasting initialization (this work)	$0.34\% \pm 0.02$
The same model without initialization	$0.63\% \pm 0.02$

Since the spatial contrasting loss is a differentiable estimation that can be computed within a network in parallel to supervised losses, future work will attempt to embed it as a semi-supervised model. This usage will allow to create models that can leverage both labeled and unlabeled data, and can be compared to similar semi-supervised models such as the ladder network Rasmus et al. (2015). It is also apparent that contrasting can occur in dimensions other than the spatial, the most straightforward is the temporal one. This suggests that similar training procedure can be applied on segments of sequences to learn useful representation without explicit supervision.

REFERENCES

- Vassileios Balntas, Edward Johns, Lilian Tang, and Krystian Mikolajczyk. Pn-net: Conjoined triple deep network for learning local image descriptors. *arXiv preprint arXiv:1601.05030*, 2016.
- Jane Bromley, James W Bentz, Léon Bottou, Isabelle Guyon, Yann LeCun, Cliff Moore, Eduard Säckinger, and Roopak Shah. Signature verification using a siamese time delay neural network. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(04):669–688, 1993.
- Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pp. 539–546. IEEE, 2005.
- Adam Coates and Andrew Y Ng. Learning feature representations with k-means. In *Neural Networks: Tricks of the Trade*, pp. 561–580. Springer, 2012.
- Adam Coates, Andrew Y Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *International Conference on Artificial Intelligence and Statistics*, pp. 215–223, 2011.
- Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1422–1430, 2015.
- Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *Advances in Neural Information Processing Systems*, pp. 766–774, 2014.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2672–2680, 2014.
- Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *International Conference on Artificial Intelligence and Statistics*, pp. 297–304, 2010.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.
- Geoffrey E Hinton. To recognize shapes, first learn to generate images. *Progress in brain research*, 165:535–547, 2007.

- Elad Hoffer and Nir Ailon. Deep metric learning using triplet network. In *Similarity-Based Pattern Recognition*, pp. 84–92. Springer, 2015.
- Ka Y Hui. Direct modeling of complex invariances for visual object features. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pp. 352–360, 2013.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of The 32nd International Conference on Machine Learning*, pp. 448–456, 2015.
- Kevin Jarrett, Koray Kavukcuoglu, Marc’Aurelio Ranzato, and Yann LeCun. What is the best multi-stage architecture for object recognition? In *Computer Vision, 2009 IEEE 12th International Conference on*, pp. 2146–2153. IEEE, 2009.
- Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Computer Science Department, University of Toronto, Tech. Rep*, 2009.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. *Advances In Neural Information Processing Systems*, pp. 1–9, 2012.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119, 2013.
- Andriy Mnih and Koray Kavukcuoglu. Learning word embeddings efficiently with noise-contrastive estimation. In *Advances in Neural Information Processing Systems*, pp. 2265–2273, 2013.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Andrew Ng. Sparse autoencoder. 2011.
- Tom Le Paine, Pooya Khorrami, Wei Han, and Thomas S Huang. An analysis of unsupervised pre-training in light of recent advances. *arXiv preprint arXiv:1412.6597*, 2014.
- Pedro O Pinheiro, Ronan Collobert, and Piotr Dollar. Learning to segment object candidates. In *Advances in Neural Information Processing Systems*, pp. 1981–1989, 2015.
- Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Antti Rasmus, Mathias Berglund, Mikko Honkala, Harri Valpola, and Tapani Raiko. Semi-supervised learning with ladder networks. In *Advances in Neural Information Processing Systems*, pp. 3532–3540, 2015.
- Ali Razavian, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. Cnn features off-the-shelf: an astounding baseline for recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 806–813, 2014.
- David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *Cognitive modeling*, 5(3):1.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 815–823, 2015.

Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–9, 2015.

Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pp. 1096–1103. ACM, 2008.

Li Wan, Matthew Zeiler, Sixin Zhang, Yann L Cun, and Rob Fergus. Regularization of neural networks using dropconnect. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pp. 1058–1066, 2013.

Shuo Yang, Ping Luo, Chen Change Loy, Kenneth W Shum, and Xiaoou Tang. Deep representation learning with target coding. 2015.

Matthew D Zeiler, Dilip Krishnan, Graham W Taylor, and Rob Fergus. Deconvolutional networks. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pp. 2528–2535. IEEE, 2010.

Junbo Zhao, Michael Mathieu, Ross Goroshin, and Yann Lecun. Stacked what-where auto-encoders. *arXiv preprint arXiv:1506.02351*, 2015.

7 APPENDIX

Table 4: Convolutional models used, based on Lin et al. (2013), Rasmus et al. (2015)

Model		
STL10	CIFAR-10	MNIST
Input: 96×96 RGB	Input: 32×32 RGB	Input: 28×28 monochrome
5×5 conv. 64 BN ReLU	3×3 conv. 96 BN LeakyReLU	5×5 conv. 32 ReLU
1×1 conv. 160 BN ReLU	3×3 conv. 96 BN LeakyReLU	
1×1 conv. 96 BN ReLU	3×3 conv. 96 BN LeakyReLU	
3×3 max-pooling, stride 2	2×2 max-pooling, stride 2 BN	2×2 max-pooling, stride 2 BN
5×5 conv. 192 BN ReLU	3×3 conv. 192 BN LeakyReLU	3×3 conv. 64 BN ReLU
1×1 conv. 192 BN ReLU	3×3 conv. 192 BN LeakyReLU	3×3 conv. 64 BN ReLU
1×1 conv. 192 BN ReLU	3×3 conv. 192 BN LeakyReLU	
3×3 max-pooling, stride 2	2×2 max-pooling, stride 2 BN	2×2 max-pooling, stride 2 BN
3×3 conv. 192 BN ReLU		
1×1 conv. 192 BN ReLU		
1×1 conv. 192 BN ReLU		
Spatial contrasting criterion		
3×3 conv. 256 ReLU	3×3 conv. 192 BN LeakyReLU	3×3 conv. 128 BN ReLU
3×3 max-pooling, stride 2	1×1 conv. 192 BN LeakyReLU	1×1 conv. 10 BN ReLU
dropout, $p = 0.5$	1×1 conv. 10 BN LeakyReLU	global average pooling
3×3 conv. 128 ReLU	global average pooling	
dropout, $p = 0.5$		
fully-connected 10		
10-way softmax		




Figure 2: First layer convolutional filters after spatial-contrasting training